L4S Architecture: Low Latency, Low Loss, Scalable Throughput Internet Service

draft-ietf-tsvwg-l4s-arch-00

B. Briscoe, K. De Schepper, M. Bagnulo

IETF-99 Jul 2017

The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

Recap

- Motivation
  - Extremely low queuing delay for all Internet traffic
- Architecture

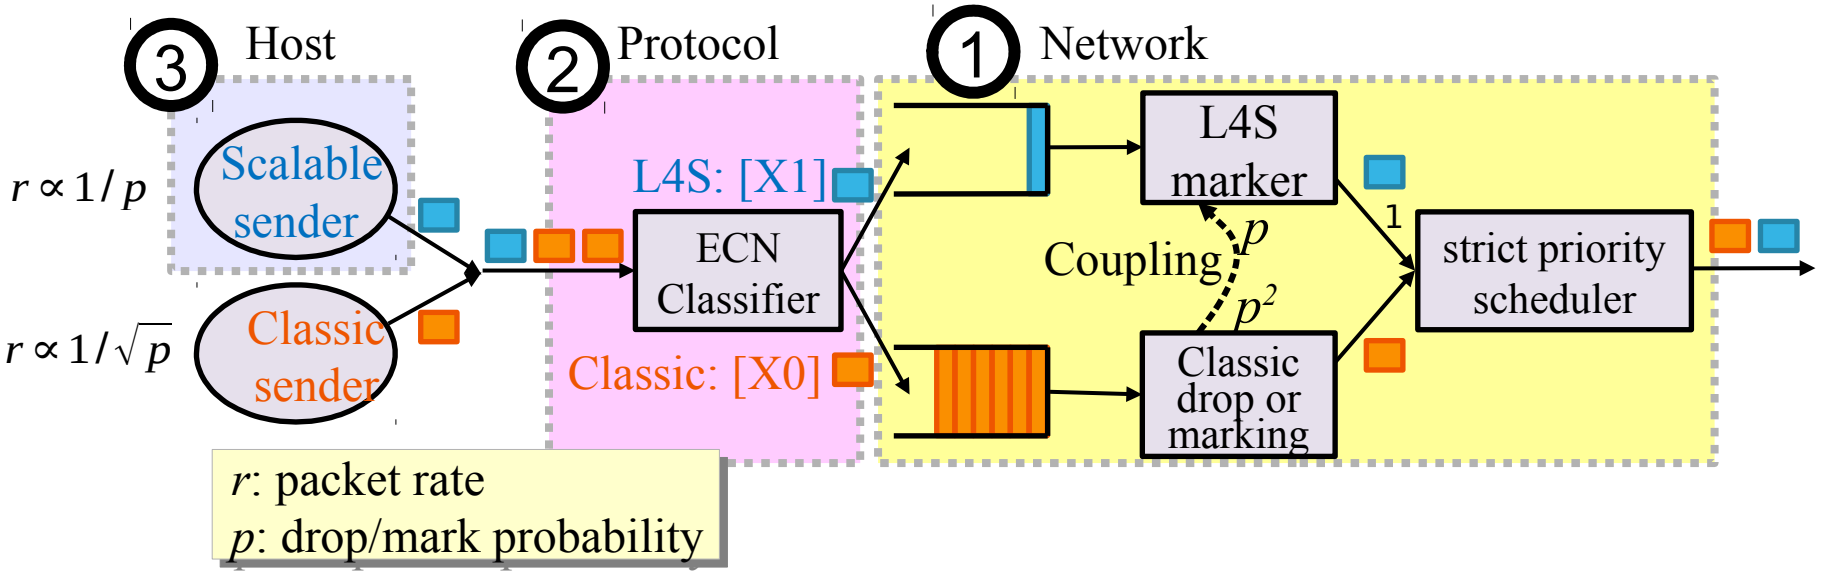

tsvwg-l4s-arch status

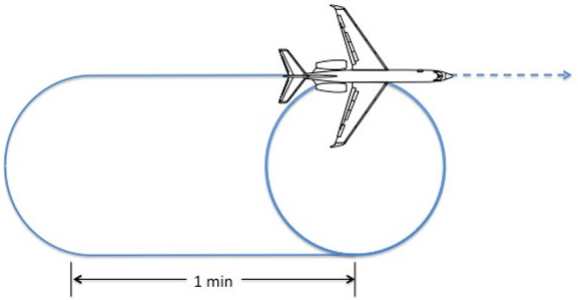

- Adopted. Now in a holding pattern
  - pending possible changes to main assumptions about the “TCP Prague Requirements”
  - Reviews, comments from implementers etc, obviously welcome
  - Same applies to tsvwg-ecn-l4s-id and tsvwg-aqm-dualq-coupled
- Two types of TCP Prague requirements:
  - Safety
  - Performance Improvements
    (e.g. rapid dynamics – presented yesterday in ICCRG)
- Come out of holding pattern when safety reqs met
  - Q-delay already 1-2 orders of magnitude better than state of the art
  - 500 µs vs 5-15 ms (FQ-CoDel or PIE)

L4S Status Update (TBA)

- Source Code
  - Dual Queue Coupled AQM, DualPI2 for Linux [UPDATED internally, release shortly]
  - Data Centre TCP (DCTCP) for Linux (in the mainline kernel), FreeBSD patch, ns2 patch.
  - Accurate ECN TCP Feedback for Linux [UPDATED, but still not fully tested]
- IETF specs
  - Low Latency, Low Loss, Scalable Throughput (L4S) Internet Service: Architecture <draft-briscoe-tsvwg-l4s-arch> [Adopted by IETF tsvwg]
  - A proposed new identifier for Low Latency, Low Loss, Scalable throughput (L4S) packets <draft-briscoe-tsvwg-ecn-l4s-id> [Adopted by IETF tsvwg]
  - enabled by <draft-ietf-tsvwg-ecn-experimentation> [Completed WGLC]
  - Dual-queue AQM: : <draft-ietf-tsvwg-aqm-dualq-coupled> [UPDATED with overload pseudocode]
  - scalable TCP algorithms, e.g. Data Centre TCP (DCTCP) <draft-ietf-tcpm-dctcp>, TCP Prague [In Progress]
  - Accurate ECN: <draft-ietf-tcpm-accurate-ecn> [UPDATED]
  - ECN++ Adding ECN to TCP control packets: <draft-ietf-tcpm-generalized-ecn> [UPDATED & Adopted]
  - ECN support in trill <draft-ietf-trill-ecn-support>, motivated by L4S [Completed WGLC]
  - ECN in QUIC <draft-johansson-quic-ecn>, motivated by L4S [UPDATED individual draft]
- 3GPP Proposal
  - ECN visibility to Radio Link Control (RLC) layer, motivated by L4S [Discussed, decision deferred to Aug'17]
- Papers
  - Article in the IETF Journal describing the Demo in Bits-N-Bites at the IETF in Prague, July 2015. “Ultra-Low Delay for All” IETF Journal, Nov 2015.
  - “Ultra-Low Delay for All: Live Experience, Live Analysis“, Proc. ACM Multimedia Systems; Demo Session (May 2016).
  - “PI2: A Linearized AQM for both Classic and Scalable TCP,” Proc. ACM CONEXT 2016 (To appear Dec 2016).
  - “Data Centre to the Home: Deployable Ultra-Low Queuing Delay for All” [Rejected – non-novel wrt IETF → Journal submission]

Identifying Modified ECN Semantics for Ultra-Low Queuing Delay

draft-ietf-tsvwg-ecn-l4s-id-00

Bob Briscoe
bobbriscoe.net

- Koen De Schepper, Inton Tsang

NOKIA Bell Labs

IETF-99 Jul 2017

The authors were part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

Premium Service vs. Default?

Codepoint	ECN bits
Not-ECT	00
ECT(0)	10
ECT(1)	01
CE	11

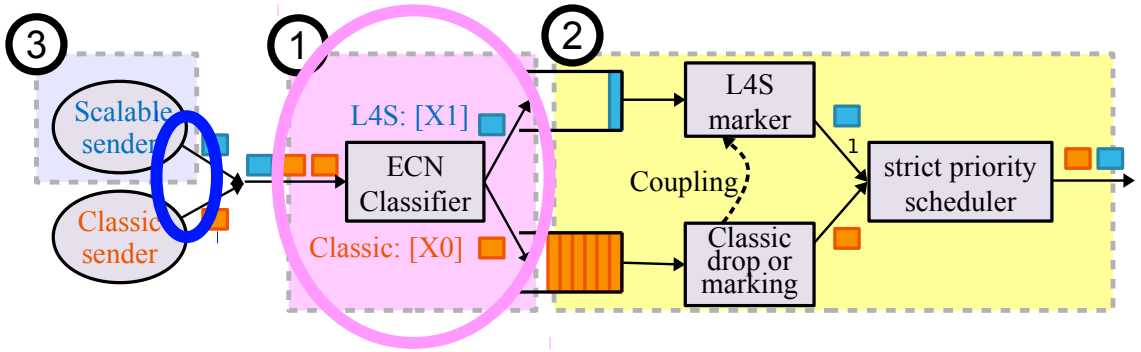

- Classifier on 2-bit ECN field in IP header (v4 or v6)
  - if ECT(1) or CE, forward to L4S
  - adopted for standardisation by IETF

AND
optionally

- Classifier on any other field
  - src/dst IP address
  - VLAN ID,
  - DSCP (local or global?)
  - bearer, …

- ECN field works end-to-end
  - network could solely enable L4S for certain addresses
  - later, could enable for all addresses
- in all cases, no packet inspection deeper than IP
  - compatibility with all privacy technology

Load balancers and ToS byte

TCP packet type	AccECN f/b	RFC3168 f/b	Congestion response
SYN	ECT	not-ECT	Reduce IW
SYN-ACK	ECT	ECT	Reduce IW
Pure ACK	ECT	ECT	None or optionally AckCC [RFC5690]
Window Probe	ECT	ECT	Usual
FIN	ECT	ECT	None or optionally AckCC [RFC5690]
RST	ECT	ECT	N/A
Re-XMT	ECT	ECT	Usual

- ECN++
  - ECN on all TCP control packets
  - draft-ietf-tcpm-generalized-ecn (tcpm on Monday)
  - all packets of one flow have consistent ECN field
  - removes need for exceptions in load balancers
- NB: still exceptions
  - ECT can change to CE
  - Legacy 8-bit classifier hardware still needs to change to 7 bits

Codepoint	ECN bits
Not-ECT	00
ECT(0)	10
ECT(1)	01
CE	11

draft-ietf-tsvwg-l4s-arch
draft-ietf-tsvwg-ecn-l4s-id

Q&A